arXiv:1607.06178v1 [cs.CV] 21 Jul 2016

# Feature Descriptors for Tracking by Detection: a Benchmark

Alessandro Pieropan Mårten Björkman Niklas Bergström Danica Kragic

***Abstract*— In this paper, we provide an extensive evaluation of the performance of local descriptors for tracking applications. Many different descriptors have been proposed in the literature for a wide range of application in computer vision such as object recognition and 3D reconstruction. More recently, due to fast key-point detectors, local image features can be used in online tracking frameworks. However, while much effort has been spent on evaluating their performance in terms of distinctiveness and robustness to image transformations, very little has been done in the contest of tracking. Our evaluation is performed in terms of distinctiveness, tracking precision and tracking speed. Our results show that binary descriptors like ORB or BRISK have comparable results to SIFT or AKAZE due to a higher number of key-points.**

## I. INTRODUCTION

Local regions of interest or key-point descriptors are widely used in Computer Vision for application such as object recognition and retrieval, 3D reconstruction and motion tracking. SIFT [1] is widely considered as one of the most robust feature descriptors, providing distinctiveness and invariance to common image transformations such as rotation and scale. However such a robustness comes at a considerable computational cost.

Recently there is much concern about efficiency caused mainly by two important factors. First there is a stable growth of portable camera enabled devices with limited computing power. Second, databases used for computer vision application are steadily increasing in size. As a consequence, there is a growing interest within the computer vision community it fast key-point detectors and binary descriptors that can dramatically decrease the computational cost of detecting and matching local regions of interest. BRIEF [2] feature descriptor in combination with FAST [3] key-point detector is among the first attempts in this direction making it suitable for real time applications. However, despite the improvement in performance, BRIEF descriptor is not very robust to image transformations. This underlines the difficulty in finding a good compromise between two competing characteristics: distinctiveness and fast computation.

More attempts has been done in this direction. BRISK [4] and ORB [5] include some modification of the BRIEF and FAST in order to achieve scale and rotation invariance. The mentioned descriptors are faster than SIFT and have comparable matching precision in case of small image transformation. More recently the binary feature descriptor KAZE [6] had comparable results to SIFT on a standard dataset [7] designed to evaluate the robustness of local descriptors on several image transformations. Moreover, by using fast non-linear filtering techniques Fast Explicit Diffusion [8], its accelerated version AKAZE [9] has shown to be competitive to SIFT in computational cost.

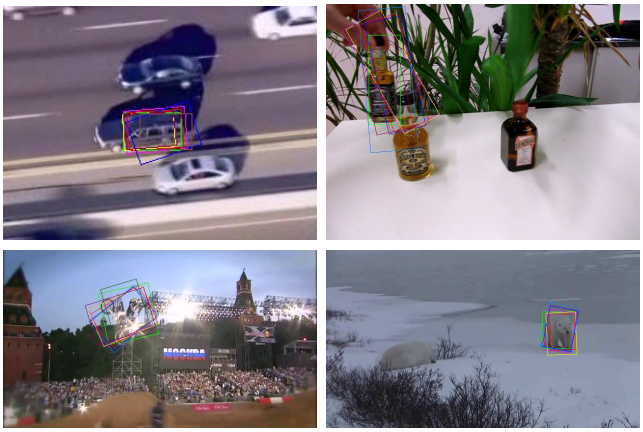

Fig. 1. Examples showing some of the videos the descriptors has been tested on and the tracking results expressed as a coloured bounding box. Clearly some feature descriptors show more precise tracking.

The main reason why AKAZE could outperform SIFT relies in the use of non-linear filtering techniques. However, given the increasing use of GPUs in computer vision, it is not clear if such techniques can be easily parallelizable on a GPU architecture as opposed to Gaussian filtering employed in SIFT. A recent work [10] has shown that it is possible to deploy AKAZE descriptor on a specialized hardware and achieve a good speedup compared to the original implementation. By implementing the descriptor in CUDA we intend to analyze its performance using a GPU and compare it against the GPU implementations of SIFT descriptor.

These fast local descriptors not only improve the performance of vision tasks such as object retrieval and 3D reconstruction but they can be used in real-time tracking system. Yet, to the best of our knowledge, very little work has been done in evaluating key-points descriptor for the specific purpose of tracking. This work aims to target that issue by providing a fair and comprehensive evaluation of the most well known descriptors. We use the recall-precision measure to evaluate the matching distinctiveness of the features and we calculate the tracking precision by integrating each local descriptor in a key-point based tracker [11]. Since there is

This research has been supported by he Japan Society for the Promotion of Science (JSPS)

The GPU used for this research was donated by the NVIDIA Corporation.

MB and DN are with CVAP/CAS, KTH, Stockholm, Sweden, celle,dani@kth.se. AP and NB are with the University of Tokyo, Japan, alessandro_pieropan, niklas_bergstrom@ipc.i.u-tokyo.ac.jp.

no dataset designed to test local descriptors for tracking purposes, the videos used in the evaluation are gathered from well known datasets used to evaluate tracking algorithms in general, described in [12]–[14]. We want to provide a practical guideline to descriptors for the specific task of tracking since it is not clear yet that a descriptor designed for image recognition is well suited for tracking.

The contribution of this work is four-fold:

**Dataset.** We collected 47 sequences from different public available datasets and annotated the ground-truth in a standardized format to ease the evaluation.

**Cuda AKAZE.** Given the growing interest in the AKAZE feature descriptor and the lack of a GPU implementation, we provide our own using CUDA.

**Evaluation Library.** We integrated the most well known descriptors in a state-of-the-art key-point based 2D tracker and we provide an interface to enable the integration of new feature descriptors.

**Evaluation.** There are three criteria the local descriptors are tested upon: distinctiveness, tracking precision and speed. First we measure the distinctiveness by matching the feature descriptors extracted in the first frame of the sequence and by computing the recall-precision. Then we calculate the tracking by integrating each feature descriptor in a key-point based tracker and we calculate the well known overlap accuracy for low, medium and high precision requirements. Last we profile the performance of each descriptor.

The dataset, our AKAZE implementation and all the code to perform the benchmark will be publicly available.

## II. Related Work

Evaluating the performance has a crucial role in computer vision in order to test the reliability of the algorithms and determine if they are of any practical use [12], [15], [16]. Evaluating feature descriptors is not an exception. One of the most relevant work has been done in 2005 by Mikolajczyk et al. [7] where affine region detectors are evaluated in terms of repatability and accuracy. SIFT feature descriptor [1] and its extension GLOH [7] had the best results on the dataset decribed in [17]. In recent years, due to a steady increase of portable devices and large scale databases, there is more attention to the runtime requirements of such descriptors. SURF [18] is among the first descriptors that presented lower computational requirements than SIFT and comparable precision. More recently, Calonder et al. [2] proposed to use a binary descriptor BRIEF, as opposed to gradient histograms used in SIFT, in combination with a fast corner detector FAST [3] achieving a fraction of the runtime performance of SURF. However, the descriptor has comparable results only when small transformations are applied to an image since the it is not scale or rotation invariant. ORB descriptor (Oriented Fast and Rotated Brief) proposed by Rublee et al. [5] enhances BRIEF adding the rotation invariance. Another binary descriptor, BRISK [4], provides both scale and orientation invariance. Its corner detector, AGAST [19] , is also faster than FAST however its invariance properties influence the overall performance.

Heinly et al. [20] evaluated exhaustively these three binary descriptors using SURF and SIFT as baseline on the Oxford dataset [17]. Not surprisingly, SIFT was the most precise while BRIEF was the fastest one but, since it is not rotation and scale invariant, its performance is significantly lower than the other descriptors upon large image transformation. This evidence underlines the challenge in finding a good compromise between precision and speed. Nevertheless, the small computational requirements of such descriptors drawn the attention on their potential use in real-time application such as tracking or simultaneous mapping and localization (SLAM).

A promising binary feature, KAZE, has been proposed by Alcantarilla et al. [6]. Its main strength relies on using non-linear diffusion filtering techniques [21] to build the scale space as opposed to Gaussian blurring. Since the latter smooths without distinction at any scale, it does not preserve object boundaries. On the other hand non-linear diffusion filtering techniques preserve edges, resulting in an increased distinctiveness of the feature descriptor. Once again, an improved precision comes at a cost in speed, however by the use of Fast Explicit Diffusion [8] the accelerated version of the feature descriptor AKAZE [9] the descriptor could achieve better accuracy than SIFT. Yet the experiments has been performed on the Oxford dataset, designed to evaluated the precision of feature descriptors with image transformation. Such features are now employed in real-time system such as tracking by detection [11], [13], [22] or more in general SLAM [23].

## III. Benchmark

The goal of the conducted benchmarking presented in detail in the experimental evaluation is three-fold. First, we want to measure the descriptiveness of the feature descriptors for matching purposes. This is crucial for the recovery of a tracker upon loss of tracking due to occlusion, changed lighting conditions, and alike. Second, we measure the tracking accuracy by integrating each feature descriptor in our 2D tracker and compute the overlap measure using the estimated object position. Third, we profile each separate step required in tracking by detection (key-point detection, descriptor computation and feature matching) in order to evaluate the performance of the feature descriptors for real-time applications.

### A. Dataset

There are several publicly available datasets designed for tracking. However, there is no agreement in the community of the standard structure to store the data. For this reason and in order to facilitate the evaluation, we collected the videos from the different datasets and standardized how these are stored. Each video is stored as a sequence of images while the ground truth, represented by an oriented bounding box of the tracked object, is saved in a separate file where each row corresponds to an image frame defined by the 8 values representing the pairs (x,y) of each vertex of the frame. The

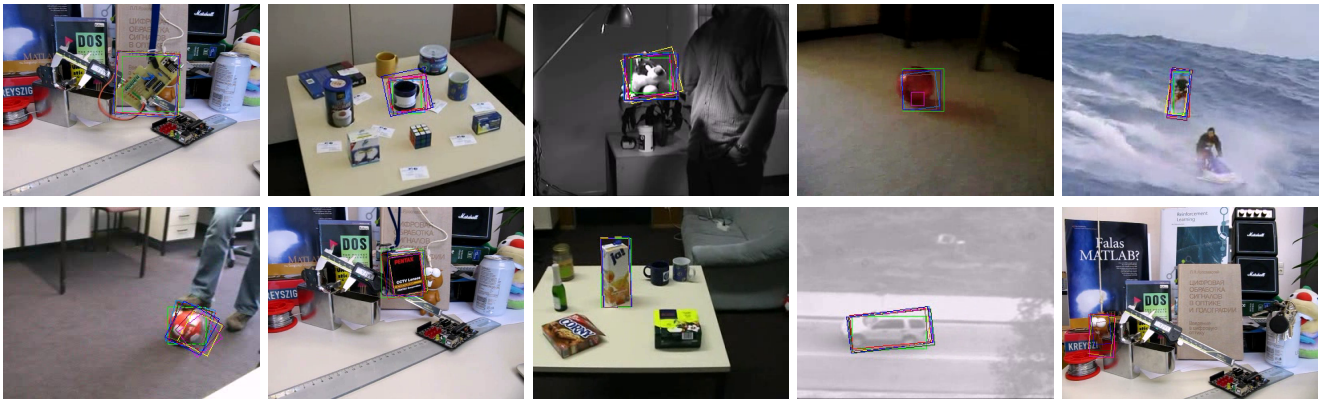

Fig. 2. A few example images from the videos used in the benchmark. The sequences include both indoor and outdoor scenes, subjected to changing lightning conditions, occlusion, and significant changes in object appearance.

generated dataset will be released publicly, along with all the code to perform the experiments and run the evaluation.

### B. Matching Evaluation Criteria

Given a sequence of images $I_1, ..., I_n$ with corresponding ground truth bounding boxes $B_1^g, ..., B_n^g$ around a target object, we extract a set of key-points $K_1$ from the first image of the sequence and regard all key-points within $B_1^g$ as descriptors of the object. For each subsequent image key-points $K_t$ are similarly extracted and then matched to those in the initial set $K_1$, generating a list of matches $M_t = \{(i,j);\ p_{1,i} \in K_1, p_{t,j} \in K_t\}$. Given the bounding boxes a key-point $p_{t,j}$, with a match $(i,j) \in M_t$, is then labeled as follows:

$$\begin{cases} \text{true positive}, & \text{if } p_{t,j} \in B_t^g \wedge p_{1,i} \in B_1^g \\ \text{false positive}, & \text{if } p_{t,j} \notin B_t^g \wedge p_{1,i} \in B_1^g \\ \text{false negative}, & \text{if } p_{t,j} \in B_t^g \wedge p_{1,i} \notin B_1^g \end{cases} \quad (1)$$

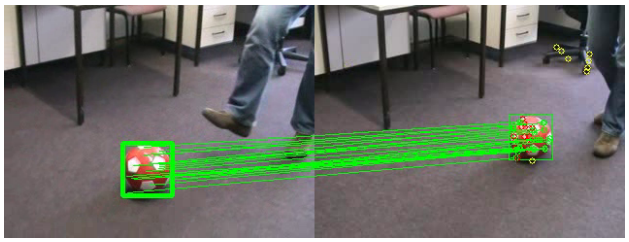

Fig. 3. Example showing how the matching precision is calculated. The image to the right shows the initial frame of the sequence. Matches in green are true positives. Circles in red are false positives. Yellow circles are false negatives, feature descriptors that are inside the object in the current frame but have a match with the background in our model.

The average ratio of true positives assesses the ability of a tracker to detect the object and potentially recover after the loss of tracking. False negatives are those feature descriptors that appear in the current frame but have no corresponding match to the initial test set $K_1$. This may commonly happen when there is a drastic change in appearance of the object. The ratio of false positives is very important to consider since it indicates the average number of outliers that will be used to estimate the pose of the object, resulting in a bad pose estimation if no additional filtering techniques are employed. A widely used filtering technique is to discard all matched key-points if the ratio between the score of the best match and the second best match is below a certain threshold $\rho$. We define a key-point as ambiguous if this criteria is not met. The number of ambiguous true and false positives is then calculated to evaluate the distinctiveness of the descriptors and evaluate the influence of this common filtering technique on the results.

### C. Evaluating tracking precision

We employ a sparse key-point based tracker to measure the precision of the feature descriptors for tracking. The tracker is initialised using a bounding box in the initial image of a tracking sequence. Inside the bounding box, feature descriptors are extracted and represent the *model* of the tracked object. The algorithm estimates the position of the object, represented as an oriented bounding box, with a combination of sparse optical flow and feature matching as outlined in Algorithm 1. For the complete details please refer to [11].

**Data**: $I_1, ..., I_n, B_1$
**Result**: $B_2, ..., B_n$
$K_1 \leftarrow$ **extract_points**$(I_1, B_1)$;
**for** $i \leftarrow 2 : n$ **do**
  $K_i^* \leftarrow$ track_points$(K_{i-1}, I_{i-1}, I_i)$;
  $B_i \leftarrow$ estimate_pose$(K_i^*)$;
  $K_i' \leftarrow$ **extract_points**$(I_i, B_i)$;
  $M \leftarrow$ **match_points**$(K_1, K_i')$;
  $K_i \leftarrow$ merge_keypoints$(K_i^*, M)$;
**end**

**Algorithm 1:** Overview of the tracking algorithm used to compute the tracking precision. The feature descriptors are employed in the steps written in bold.

Our original tracking algorithm used ORB (Oriented Fast and Rotated Brief) proposed by Rublee et al. [5]. We extended the tracking algorithm making it more modular so that tracking can be performed with various feature descriptors

that we evaluate. To estimate the tracking precision, we used the widely accepted overlap measure:

$$\Theta(B_t, B_t^g) = \frac{B_t \cap B_t^g}{B_t \cup B_t^g} \quad (2)$$

where $B_t$ is the bounding box estimated by our tracker and $B_t^g$ is the bounding box provided by the ground truth. We define three precision requirements $\Upsilon \in \{0.25, 0.5, 0.75\}$ that indicate low, medium and high tracking accuracy. This is a more indicative evaluation compared to the overall accuracy. For instance, an overall value of 0.5 is ambiguous because it may indicate either a stable average accuracy around the value or a very precise evaluation in part of the video while poor in the rest. The overall accuracy of the tracker is then computed for each precision requirement $\Upsilon$ as the fraction of all generated bounding boxes for which $\Theta(B_t, B_t^g) > \Upsilon$.

### D. Parameter setting

All the employed descriptors require many parameters to be initialised and run properly. In order to achieve the most fair comparison between them, we decided to keep the values suggested in the original publication of the feature descriptor or the implementation if such was available. However, there are some parameters that most of the feature descriptors share and that we set to the same value. Thus, the maximum number of features extracted is set to 2500 and the number of scale space octaves is set to 4.

## IV. Results

In this section, we summarise the results of our assessment. The descriptors considered in the evaluation are: BRISK, ORB, SURF and SIFT included in OpenCV library. The implementation of AKAZE made available by its original author [9] and the implementation of SIFT using CUDA presented in [24].

### A. Distinctiveness

The first aspect was to assess the distinctiveness. Table I shows the average number of key-points extracted, descriptors belonging to the object to track and the ratio of true positives (TP), false positives (FP) and true positives that pass the second best match ratio check (TTP).

TABLE I

Average number of feature extracted, object features, true positives and false positives. Every row is normalized by its maximum value.

| | AKAZE | BRISK | AKAZE(C) | SIFT(C) | ORB | SIFT | SURF |
|---|---|---|---|---|---|---|---|
| points | 0.91 | 0.82 | 0.53 | 0.42 | 1.0 | 0.99 | 0.34 |
| object | 0.66 | 0.67 | 0.51 | 0.32 | 1.0 | 0.63 | 0.23 |
| TP | 0.84 | 0.88 | 0.95 | 1.0 | 0.92 | 0.72 | 0.76 |
| FP | 0.94 | 0.92 | 0.89 | 0.86 | 0.9 | 1.0 | 0.98 |
| TTP | 0.72 | 0.59 | 0.71 | 1.0 | 0.53 | 0.55 | 0.71 |

It is interesting to notice that BRISK, ORB and SIFT extract a higher number of feature descriptors in general. In particular, BRISK and ORB have a higher number of key-points extracted within the area of the object. However, looking at the average amount of true positives, it can be seen that the best performing descriptors are AKAZE and the implementation of SIFT on the GPU. This is a first indicator of the quality of the descriptors extracted. Moreover, it can be noticed that the true positives are also more distinctive in the case of AKAZE and SIFT since the number of TTP is higher.

### B. Tracking accuracy

As explained in the previous section, we evaluated the accuracy of the feature descriptors by running our tracker and calculating the overlap measure for low, medium and high accuracy requirements.

Fig. 4. Examples showing the behaviour of the feature descriptors upon occlusion. Upon recovery from track loss more descriptive descriptors allow the tracker to recover faster.

Table II summarizes tracking results on all the video sequences included in the dataset. Our experiments show that AKAZE, BRISK, ORB and SIFT have comparable results. It is interesting to note that BRISK and ORB compensate their weak distinctiveness with a higher amount of weak descriptors extracted. A high number of feature points proved to be effective in tracking in the video sequences where the object suffers drastic scale changes and full occlusion, making the recovery after track loss faster, see Fig. 4. We also noticed that AKAZE, more than SIFT, suffers the change in scale.

### C. Tracking performance

The dataset used for benchmarking includes video sequences of various resolution. Fig. 5 shows the average performance of each feature descriptor on the resolutions having the highest number of sequences.

The two most important factors that influence performance are resolution and number of key points extracted. The former influences particularly the detection step when the scale space of descriptors is computed and key-points are detected. The latter influences more the extraction step when feature descriptors are calculated and the matching step. The average performance of each separate step can be seen in Fig. 6. One interesting aspect to notice is the variance of the performance in Fig. 5: BRISK, ORB and SIFT have the higher variance while CUDA SIFT and CUDA AKAZE have lower variance. This is also a good indicator of the level of parallelism of the implementation of the descriptor. It is interesting to notice that the computation of the non linear scale space required by AKAZE is not perfectly suited for a

TABLE II

TRACKING RESULTS WITH LOW, MEDIUM AND HIGH ACCURACY REQUIREMENTS. THE HIGH NUMBER OF KEY POINTS EXTRACTED BY ORB OR BRISK COMPENSATE THEIR WEAK DESCRIPTORS. THIS COMES WITH A COST IN PERFORMANCE.

| feature | AKAZE | BRISK | AKAZE(C) | SIFT(C) | ORB | SIFT | SURF |
|---|---|---|---|---|---|---|---|
| ball | 0.94/**0.86**/0.18 | **0.99**/0.53/0.15 | 0.97/0.76/0.18 | 0.95/0.64/0.16 | 0.96/0.4/0.12 | 0.95/0.51/0.11 | 0.83/0.65/**0.19** |
| basketball | **0.21**/**0.09**/**0.03** | 0.03/0.03/0.02 | 0.1/**0.09**/**0.03** | 0.05/0.05/0.01 | 0.03/0.02/0.01 | 0.03/0.03/0.01 | 0.03/0.02/0.02 |
| bicycle | 0.2/0.17/0.07 | **0.64**/**0.61**/**0.33** | 0.5/0.43/0.21 | **0.64**/**0.61**/0.28 | **0.64**/0.57/0.26 | **0.64**/0.6/0.3 | 0.3/0.28/0.11 |
| board | 0.88/**0.87**/0.16 | **0.95**/0.86/0.22 | 0.89/0.86/0.19 | 0.93/0.86/0.18 | 0.85/0.83/0.18 | 0.87/0.83/**0.23** | 0.87/0.86/0.21 |
| bolt | **0.03**/**0.01**/**0.0** | **0.03**/**0.01**/**0.0** | **0.03**/**0.01**/**0.0** | 0.02/**0.01**/**0.0** | **0.03**/**0.01**/**0.0** | **0.03**/**0.01**/**0.0** | **0.03**/**0.1**/**0.0** |
| box | 0.63/0.55/0.34 | 0.95/0.89/**0.72** | 0.94/0.9/**0.72** | **0.96**/**0.92**/0.62 | 0.91/0.87/0.69 | 0.87/0.82/0.55 | 0.83/0.8/0.57 |
| car | 0.07/0.07/0.06 | **0.64**/**0.63**/0.32 | 0.0/0.0/0.0 | **0.64**/**0.63**/0.51 | **0.64**/**0.63**/**0.59** | **0.64**/**0.63**/0.2 | 0.52/0.35/0.12 |
| car 2 | 0.86/0.85/0.56 | 0.86/0.85/0.58 | 0.86/0.86/0.63 | 0.65/0.65/0.52 | **0.87**/0.86/0.63 | 0.86/0.86/0.63 | 0.87/**0.87**/**0.65** |
| carchase | 0.18/0.11/0.04 | 0.4/0.36/0.14 | 0.24/0.2/0.06 | 0.4/**0.37**/0.12 | **0.41**/0.37/**0.16** | 0.39/0.36/0.15 | 0.14/0.14/0.08 |
| cup | 0.96/0.94/0.59 | **1.0**/**1.0**/0.68 | **1.0**/**1.0**/**0.72** | **1.0**/**1.0**/0.66 | **1.0**/**1.0**/0.69 | 0.94/0.93/0.65 | 0.65/0.65/0.64 |
| david | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** |
| diving | 0.14/0.13/**0.08** | **0.28**/0.13/0.07 | 0.17/0.13/**0.08** | 0.16/0.13/0.07 | 0.19/0.13/0.07 | 0.18/**0.14**/**0.08** | 0.18/**0.14**/**0.08** |
| dog | **1.0**/**1.0**/0.79 | **1.0**/**1.0**/0.84 | **1.0**/**1.0**/0.86 | **1.0**/**1.0**/0.83 | **1.0**/**1.0**/**0.87** | 0.96/0.95/0.81 | 1.0/0.99/0.72 |
| drunk | 0.17/**0.16**/0.03 | 0.04/0.04/0.02 | 0.18/0.1/0.03 | 0.11/0.03/0.01 | **0.2**/0.1/0.02 | 0.11/0.1/**0.06** | 0.08/0.08/0.04 |
| fernando | 0.18/0.11/**0.02** | 0.19/0.12/**0.02** | **0.2**/0.11/**0.02** | 0.18/0.05/**0.02** | 0.19/**0.14**/**0.02** | 0.17/0.08/**0.02** | 0.14/0.05/**0.02** |
| fish1 | 0.04/0.01/**0.0** | 0.05/0.01/**0.0** | 0.04/0.01/**0.0** | 0.05/**0.02**/**0.0** | **0.06**/0.01/**0.0** | **0.06**/0.01/**0.0** | 0.05/0.01/**0.0** |
| fish2 | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** | **0.0**/**0.0**/**0.0** |
| gym | **1.0**/**0.98**/**0.24** | **1.0**/0.89/0.14 | **1.0**/0.97/0.17 | 0.58/0.51/0.1 | **1.0**/0.69/0.15 | 0.96/0.77/0.13 | 0.29/0.26/0.09 |
| gymnastics | 0.44/0.4/**0.36** | 0.45/0.42/0.35 | 0.45/0.42/**0.36** | 0.44/0.42/**0.36** | 0.45/0.42/**0.36** | **0.55**/0.43/0.35 | 0.46/**0.44**/0.35 |
| hand1 | 0.33/**0.24**/**0.04** | 0.19/0.08/0.03 | **0.35**/0.21/**0.04** | 0.17/0.1/0.03 | 0.22/0.12/0.03 | 0.19/0.09/**0.04** | 0.1/0.09/**0.04** |
| hand2 | 0.16/0.06/0.04 | 0.2/0.08/0.04 | 0.13/0.06/**0.05** | 0.11/0.06/0.04 | **0.23**/**0.09**/0.04 | 0.16/0.06/0.04 | 0.11/0.06/0.04 |
| jogging | **0.95**/**0.87**/0.2 | 0.64/0.61/**0.22** | 0.89/0.73/0.13 | 0.67/0.51/0.16 | 0.75/0.51/0.12 | 0.25/0.24/0.09 | 0.51/0.43/0.07 |
| juice | **1.0**/**1.0**/0.96 | **1.0**/**1.0**/**1.0** | **1.0**/**1.0**/**1.0** | **1.0**/**1.0**/**1.0** | **1.0**/**1.0**/**1.0** | **1.0**/**1.0**/**1.0** | **1.0**/**1.0**/**1.0** |
| jumping | 0.39/0.33/0.12 | 0.99/0.94/**0.5** | 0.43/0.4/0.12 | 0.93/0.9/0.5 | **1.0**/**0.95**/0.39 | 0.49/0.41/0.17 | 0.11/0.11/0.1 |
| lemming | 0.07/0.05/0.02 | 0.6/0.58/0.3 | 0.12/0.07/0.02 | 0.49/0.48/0.3 | **0.66**/**0.62**/**0.32** | 0.55/0.52/0.27 | 0.13/0.12/0.08 |
| liquor | 0.87/0.86/0.71 | **0.94**/**0.93**/0.75 | 0.89/0.89/0.73 | 0.93/0.92/**0.76** | 0.9/0.89/0.74 | 0.91/0.9/0.74 | 0.88/0.87/0.73 |
| motocross | 0.23/0.21/**0.0** | 0.37/0.18/**0.0** | 0.39/0.21/**0.0** | 0.42/0.2/**0.0** | **0.54**/**0.31**/**0.0** | 0.42/0.23/**0.0** | 0.29/0.21/**0.0** |
| mtbk | 0.59/0.57/0.26 | **1.0**/**1.0**/0.29 | **1.0**/**1.0**/0.34 | **1.0**/**1.0**/0.35 | **1.0**/**1.0**/0.35 | **1.0**/**1.0**/**0.44** | **1.0**/**1.0**/0.31 |
| person | 0.99/0.92/**0.58** | 0.94/0.82/0.45 | **1.0**/0.92/0.54 | 0.98/0.91/0.48 | 0.92/0.76/0.41 | 0.98/**0.93**/0.47 | 0.85/0.81/0.51 |
| crossing | **0.87**/0.8/0.68 | 0.79/0.75/0.65 | **0.87**/**0.81**/**0.71** | 0.84/0.79/0.7 | 0.76/0.68/0.51 | 0.83/0.73/0.59 | 0.78/0.73/0.61 |
| occluded | **1.0**/0.95/**0.82** | **1.0**/0.95/**0.82** | **1.0**/**0.97**/**0.82** | **1.0**/0.95/**0.82** | **1.0**/0.95/**0.82** | **1.0**/0.95/0.81 | **1.0**/0.95/**0.82** |
| polarbear | 0.45/**0.44**/**0.24** | 0.44/0.41/0.17 | **0.49**/0.42/0.16 | 0.0/0.0/0.0 | 0.45/0.44/0.23 | 0.44/0.38/0.12 | 0.0/0.0/0.0 |
| singer | 0.7/0.68/0.34 | **1.0**/0.83/0.33 | **1.0**/**0.99**/0.35 | **1.0**/0.94/0.31 | **1.0**/0.9/**0.39** | **1.0**/0.9/0.34 | 0.59/0.58/0.37 |
| skating | 0.1/0.1/0.02 | **0.7**/0.26/0.03 | 0.19/0.16/0.02 | 0.3/0.27/0.03 | 0.56/**0.41**/**0.04** | 0.27/0.19/0.02 | 0.12/0.1/0.02 |
| sphere | **1.0**/**1.0**/0.28 | 0.18/0.18/0.13 | **1.0**/**1.0**/**0.51** | 0.21/0.21/0.2 | 0.6/0.27/0.2 | 0.2/0.18/0.15 | 0.18/0.18/0.15 |
| sunshade | 0.86/0.57/0.09 | 0.92/0.43/0.16 | 0.71/0.26/0.09 | 0.94/0.4/0.2 | 0.88/0.41/0.15 | **0.96**/**0.63**/**0.22** | 0.49/0.19/0.05 |
| surfing | 0.88/0.79/0.45 | **1.0**/0.92/0.56 | 0.64/0.64/0.34 | **1.0**/0.93/**0.57** | 1.0/**0.99**/0.46 | **1.0**/0.97/0.27 | 0.04/0.04/0.04 |
| sylvester | 0.93/0.77/0.37 | 0.97/0.92/**0.54** | 0.92/0.83/0.39 | 0.96/**0.95**/0.52 | **1.0**/0.88/0.43 | 0.99/0.92/0.43 | 0.64/0.62/0.41 |
| torus | 0.33/**0.21**/**0.06** | 0.26/0.11/0.04 | **0.34**/0.1/0.05 | 0.18/0.06/0.05 | 0.32/0.06/0.05 | 0.25/0.06/0.05 | 0.09/0.06/0.05 |
| track | 0.38/0.38/0.25 | **0.79**/0.77/0.61 | **0.79**/**0.79**/0.67 | **0.79**/**0.79**/**0.74** | 0.75/0.74/0.54 | **0.79**/**0.79**/0.69 | 0.29/0.29/0.25 |
| trellis | 0.17/0.16/0.12 | 0.19/0.16/0.05 | 0.15/0.13/0.11 | 0.11/0.08/0.02 | 0.16/0.12/0.04 | 0.17/0.15/0.04 | **0.2**/**0.2**/**0.14** |
| tunnel | 0.16/0.11/0.03 | 0.18/0.09/0.02 | **0.32**/**0.18**/0.02 | 0.11/0.1/0.04 | 0.2/0.13/0.02 | 0.18/0.12/0.02 | 0.14/0.11/**0.05** |
| woman | 0.23/0.19/0.12 | 0.19/0.11/0.06 | 0.16/0.14/0.09 | 0.34/**0.25**/**0.13** | **0.4**/0.23/0.06 | 0.15/0.12/0.07 | 0.11/0.09/0.04 |
| overall | 0.5/0.46/0.24 | 0.58/**0.5**/**0.29** | 0.54/0.48/0.27 | 0.54/0.48/**0.29** | **0.6**/**0.5**/0.28 | 0.54/0.48/0.26 | 0.39/0.36/0.23 |

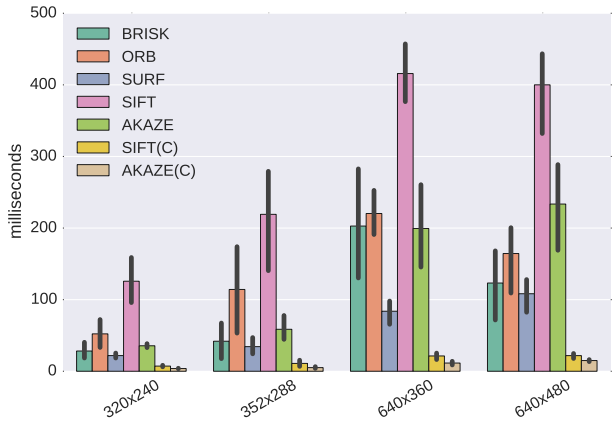


Fig. 5. Average time spent on tracking the object in a single frame: important factors are the resolution and the number of feature descriptors extracted. The variance of the results is a good indication of how much the number of feature descriptors influences the performance. It can be seen that the implementations have a lower variation due to the high level of parallelism.

GPU architecture since it requires many sequential steps, as a result the key-detector is slower that SIFT. However since the descriptor is binary, its extraction and matching compensate in terms of performance. It is also important to remember that CPU implementations do not exploit the same number of cores, complicating a fair comparison between methods.

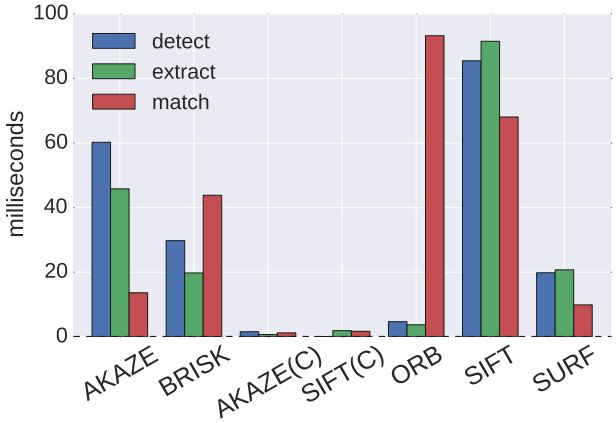


Fig. 6. Performance of the detect, extract and match steps of each feature descriptor.

### D. Discussion

Most of the feature descriptors proved to be effective for tracking purposes showing a good precision and performance. This is positive since it makes these suitable for real-time applications. Despite being assessed as somewhat weak in terms of distinctiveness, ORB and BRISK have been proven to be good for real time frameworks. However, the trade-off between the distinctiveness and accuracy is not easy to define. AKAZE and SIFT have been proven to be more distinctive as descriptors but only their implementations on the GPU allow real time performances. The most important factors that diversify the results are the characteristics of the video sequences: object transformations, in particular scale, object appearance, light condition and motion blur. AKAZE and SIFT have shown to be more distinctive and effective when the video does not present drastic movements. AKAZE seems to be particularly sensitive to change in scale or blurring, on the other hand it has higher performance on low textured objects and people, still in sequences where the tracked object keep a constant distance from the camera as in Fig. 8. Moreover, we noticed that the matching rate of all features drop consistently upon fast movements of the camera as we discussed in [11]. However, the performance of weak feature descriptors such as BRISK or ORB seems to be less sensitive to this kind of noise as in Fig. 8.

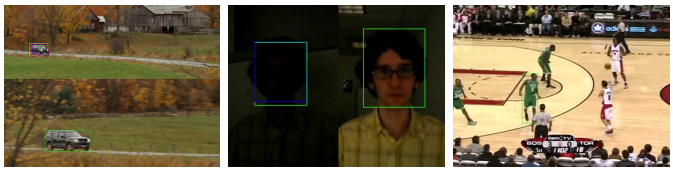

(a) scale (b) light (c) multi-instance

Fig. 7. Examples showing the main problems that feature descriptors cannot address.

There are still some issue that need to be addressed in relation to achieving a robust tracking system. First, even if many descriptors are invariant to scale or rotation, we noticed that this does not hold for drastic changes like in Fig. 7(a). One possible solution is to extract features from appearances of the object generated through synthetic transformations. It has been shown by Morel [25] that this technique improves the matching performance of SIFT descriptor. Second, feature descriptors are sensitive to light conditions, see Fig. 7(b). This is particularly relevant for robotics applications since the interaction of a robotic platform with a target object may occlude the light source. Third, the common matching approach to detect an object or compute the transformation between images [7] does not work in the presence of multiple targets with similar appearance. This is the case shown in Fig.7(c) where more players have the same outfit.

## V. Conclusion

We performed an evaluation of the most common feature descriptors for the purpose of tracking by detection. Our experiments have shown that most of the feature descriptors have comparable results on a dataset presenting challenges like motion blur, occlusion, scale, rotation and light changes. AKAZE and SIFT have proven to be more distinctive with small object transformation and motion blur. AKAZE seems to be particularly sensitive to scale changes. On the other hand ORB and BRISK have better performance on sequences with significant motion blur. Given the growing interested in AKAZE descriptor, we implemented the method using CUDA and evaluated the impact of a non-linear filtering techniques, not particularly suited for a GPU's architecture, on the performance of the AKAZE feature descriptor. Our results have shown that benefit of the extracting and matching a binary descriptor compensates the more demanding detection

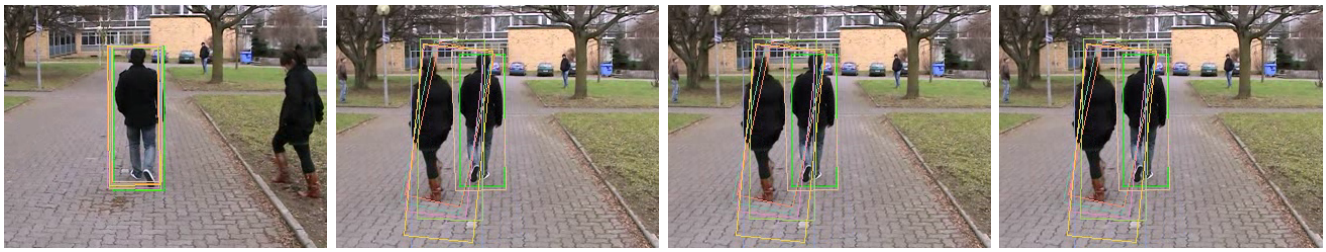

(a) Crossing: constant distance from camera, target often occluded.

(b) Jumping: high motion blur.

(c) Skating: motion blur, scale and light change.

Fig. 8. A few example where the precision of the descriptors differs consistently. In sequence (a) AKAZE the camera keeps the a constant distance from the target who is often occluded by other people. SIFT and AKAZE performs better in this scenario. In sequence (b) the target is often blurred due to high motion, ORB and BRISK have the best precision in such as scenario. Sequence (c) is one of the most challenging, it contains scale and light changes and motion blur. BRISK and ORB have higher accuracy however on light changes all the descriptors lose the target.

step achieving similar performances of the implementation of SIFT on the GPU. All the code to perform the benchmark, the dataset and our implementation of AKAZE will be publicly available in order to ease researches in this area.

## REFERENCES


[1] G. D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, pages 91–110, November 2004.

[2] C. Strecha M. Calonder, V. Lepetit and P. Fua. Brief: Binary robust independent elementary features. In *ECCV*, 2010.

[3] E. Rosten and T. Drummond. Machine learning for high-speed corner detection. In *ECCV*, 2006.

[4] S. Leutenegger, M. Chli, and R. Y. Siegwart. Brisk: Binary robust invariant scalable keypoints. In *ICCV*, 2011.

[5] K. Konolige E. Rublee, V. Rabaud and G. Bradski. Orb: An efficient alternative to sift or surf. In *ICCV*, 2011.

[6] A. Bartoli P. F. Alcantarilla and A. J. Davison. Kaze features. In *ECCV*, 2012.

[7] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *PAMI*, 2005.

[8] M. Goesele, S. Roth, A. Kuijper, B. Schiele, K. Schindler, S. Grewenig, J. Weickert, and A. Bruhn. From box filtering to fast explicit diffusion. *Pattern Recognition*, 2010.

[9] J. Nuevo P. F. Alcantarilla and A. Bartoli. Fast explicit diffusion for accelerated features in nonlinear scale spaces. In *BMVC*, 2013.

[10] G. Jiang, L. Liu, W. Zhu, S. Yin, and S. Wei. A 127 fps in full hd accelerator based on optimized akaze with efficiency and effectiveness for image feature extraction. In *Design Automation Conference*. ACM, 2015.

[11] A. Pieropan, N. Bergström, H. Kjellström, and M. Ishikawa. Robust and adaptive keypoint based object tracking. *Advanced Robotics*, 2015. In press.

[12] Y. Wu, J. Lim, and M. H. Yang. Online Object Tracking: A Benchmark. *CVPR*, 2013.

[13] G. Nebehay and R. Pflugfelder. Consensus-based matching and tracking of keypoints for object tracking. In *WCACV*, 2014.

[14] S.Hare, A. Saffari, and P. Torr. Struck: Structured output tracking with kernels. *ICCV*, 2011.

[15] H. I. Christensen and J. Phillips, editors. *Empirical Evaluation Methods in Computer Vision*. Machine Perception and Artificial Intelligence. World Scientific Press, July 2002.

[16] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012.

[17] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *IJCV*, 2005.

[18] T. Tuytelaars H. Bay, A. Ess and L. Van Gool. Speeded-up robust features (surf). *CVIU*, 2008.

[19] E. Mair, G. D. Hager, D. Burschka, M. Suppa, and G. Hirzinger. Adaptive and generic corner detection based on the accelerated segment test. *ECCV*, 2010.

[20] J. Heinly, E. Dunn, and J.M. Frahm. Comparative evaluation of binary features. In *ECCV*, 2012.

[21] J. Weickert, B. M. Ter Haar Romeny, and M. A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 1998.

[22] A. Pieropan, N. Bergström, M. Ishikawa, and H. Kjellström. Robust 3d tracking of unknown objects. In *International Conference on Robotics and Automation*, 2015.

[23] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos. ORB-SLAM: a Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics*, 2015.

[24] M. Björkman, N. Bergström, and D. Kragic. Detecting, segmenting and tracking unknown objects using multi-label {MRF} inference. *Computer Vision and Image Understanding*, 2014.

[25] J. M. Morel and G. Yu. Asift: A new framework for fully affine invariant image comparison. *SIAM Journal on Imaging Sciences*, 2009.